What Is Claimed Is:

1. A method for transparently interconnecting multiple network links into a single virtual network link, comprising:

receiving a packet at a routing bridge (Rbridge), wherein the Rbridge belongs to a set of one or more Rbridges that transparently interconnect the multiple network links into the single virtual network link;

wherein the Rbridges automatically obtain information specifying which endnodes are located on the multiple network links without the endnodes having to proactively announce their presence to the Rbridges; and

if a destination for the packet resides on the same virtual network link, routing the packet to the destination, wherein the route can be an optimal path to the destination, and is not constrained to lie along a spanning tree through the set of Rbridges.

2. The method of claim 1,

wherein endnodes on the multiple network links send packets to each other through the Internet Protocol (IP); and

wherein all endnodes on the multiple network links that comprise the single virtual network link have an address with the same prefix, or have an address with a prefix from the same set of prefixes.

3. The method of claim 1, wherein automatically obtaining the information specifying which endnodes are located on the multiple network links involves an Rbridge performing a distributed query to other Rbridges, asking them to initiate an ARP query (IPv4), or to send a neighbor discovery solicitation message (IPv6).

4. The method of claim 3, wherein an Rbridge replies to an ARP query (IPv4) or a neighbor solicitation message (IPv6) from a node on its local link regarding a target node on a different link within the virtual network link with the layer 2 address of the target node.

5. The method of claim 1, wherein if the layer 3 destination for the packet does not reside on the same virtual network link, the method further comprises routing the packet based on the layer 2 destination address specified by the source of the packet.

6. The method of claim 1, wherein routing the packet to the destination involves ensuring that the packet includes a hop count to avoid problems with temporary loops.

7. The method of claim 6, wherein the hop count is maintained in a layer 3 header of the packet.

8. The method of claim 6, wherein the hop count is added in an encapsulation header.

9. The method of claim 1, wherein the Rbridges automatically obtain information specifying where endnodes are located by examining data packets.

10. The method of claim 9,

wherein if the Rbridge receives the packet directly from a sender and the Rbridge is forwarding the packet to a downstream Rbridge, the Rbridge modifies the packet to indicate that the packet is in-transit; and

wherein if the Rbridge receives the packet from an upstream Rbridge and if the Rbridge is forwarding the packet directly to the destination, the Rbridge undoes the modification to the packet to indicate that the packet is no longer in transit.

11. The method of claim 1, wherein the Rbridges automatically obtain information specifying where endnodes are located by examining control packets.

12. The method of claim 11, wherein the control packets are IPv4 address resolution protocol (ARP) packets.

13. The method of claim 11, wherein the control packets are IPv6 neighborhood discovery packets.

14. The method of claim 1, wherein the Rbridges automatically obtain information specifying where endnodes are located by periodically polling the endnodes.

15. The method of claim 1, wherein if the Rbridge cannot determine which link the destination is located on, the packet is sent along a spanning tree through the Rbridges.

16. The method of claim 1, wherein when forwarding an IP packet to the destination link, the method uses a special source address in the layer 2 header

to indicate to Rbridges that this packet did not originate on that link, and instead was forwarded onto that link by an Rbridge.

17. A computer-readable storage medium storing instructions that when executed by a computer cause the computer to perform a method for transparently interconnecting multiple network links into a single virtual network link, the method comprising:

receiving a packet at a routing bridge (Rbridge), wherein the Rbridge belongs to a set of one or more Rbridges that transparently interconnect the multiple network links into the single virtual network link;

if a destination for the packet resides on the same virtual network link, routing the packet to the destination, wherein the route can be an optimal path to the destination, and is not constrained to lie along a spanning tree through the set of Rbridges; and

automatically obtaining information at the Rbridges specifying which endnodes are located on the multiple network links without the endnodes having to proactively announce their presence to the Rbridges.

18. The computer-readable storage medium of claim 17,

wherein endnodes on the multiple network links send packets to each other through the Internet Protocol (IP); and

wherein all endnodes on the multiple network links that comprise the single virtual network link have an address with the same prefix, or have an address with a prefix from the same set of prefixes.

19. The computer-readable storage medium of claim 17, wherein automatically obtaining the information specifying which endnodes are located on

the multiple network links involves an Rbridge performing a distributed query to other Rbridges, asking them to initiate an ARP query (IPv4), or to send a neighbor discovery solicitation message (IPv6).

20. The computer-readable storage medium of claim 17, wherein an Rbridge replies to an ARP query (IPv4) or a neighbor solicitation message (IPv6) from a node on its local link regarding a target node on a different link within the virtual network link with the layer 2 address of the target node.

21. The computer-readable storage medium of claim 17, wherein if the layer 3 destination for the packet does not reside on the same virtual network link, the method further comprises routing the packet based on the layer 2 destination address specified by the source of the packet.

22. The computer-readable storage medium of claim 17, wherein routing the packet to the destination involves ensuring that the packet includes a hop count to avoid problems with temporary loops.

23. The computer-readable storage medium of claim 22, wherein the hop count is maintained in a layer 3 header of the packet.

24. The computer-readable storage medium of claim 22, wherein the hop count is added in an encapsulation header.

25. The computer-readable storage medium of claim 22, wherein the Rbridges automatically obtain information specifying where endnodes are located by examining data packets.

26. The computer-readable storage medium of claim 25,

wherein if the Rbridge receives the packet directly from a sender and the Rbridge is forwarding the packet to a downstream Rbridge, the Rbridge modifies the packet to indicate that the packet is in-transit; and

wherein if the Rbridge receives the packet from an upstream Rbridge and if the Rbridge is forwarding the packet directly to the destination, the Rbridge undoes the modification to the packet to indicate that the packet is no longer in transit.

27. The computer-readable storage medium of claim 17, wherein the Rbridges automatically obtain information specifying where endnodes are located by examining control packets.

28. The computer-readable storage medium of claim 27, wherein the control packets are IPv4 address resolution protocol (ARP) packets.

29. The method of claim 27, wherein the control packets are IPv6 neighborhood discovery packets.

30. The computer-readable storage medium of claim 17, wherein the Rbridges automatically obtain information specifying where endnodes are located by periodically polling the endnodes.

31. The computer-readable storage medium of claim 17, wherein if the Rbridge cannot determine which link the destination is located on, the packet is sent along a spanning tree through the Rbridges.

32. The apparatus of claim 17, wherein when forwarding an IP packet to the destination link, the method uses a special source address in the layer 2 header to indicate to Rbridges that this packet did not originate on that link, and instead was forwarded onto that link by an Rbridge.

33. An apparatus that transparently interconnects multiple network links into a single virtual network link, comprising:

a routing bridge (Rbridge) that belongs to a set of one or more Rbridges that transparently interconnect the multiple network links into the single virtual network link;

a receiving mechanism within the Rbridge configured to receive a packet;

a link state determination mechanism, which is configured to automatically obtain information specifying which endnodes are located on the multiple network links without the endnodes having to proactively announce their presence to the Rbridges; and

a routing mechanism within the Rbridge, wherein if a destination for the packet resides on the same virtual network link, the routing mechanism is configured to route the packet to the destination, wherein the route can be an optimal path to the destination, and is not constrained to lie along a spanning tree through the set of Rbridges.

34. The apparatus of claim 33,

wherein endnodes on the multiple network links send packets to each other through the Internet Protocol (IP); and

wherein all endnodes on the multiple network links that comprise the single virtual network link have an address with the same prefix, or have an address with a prefix from the same set of prefixes.

35. The apparatus of claim 33, wherein while automatically obtaining the information specifying which endnodes are located on the multiple network links, the link state determination mechanism is configured to cause an Rbridge to perform a distributed query to other Rbridges, asking them to initiate an ARP query (IPv4), or to send a neighbor discovery solicitation message (IPv6).

36. The apparatus of claim 35, wherein an Rbridge replies to an ARP query (IPv4) or a neighbor solicitation message (IPv6) from a node on its local link regarding a target node on a different link within the virtual network link with the layer 2 address of the target node.

37. The apparatus of claim 33, wherein if the layer 3 destination for the packet does not reside on the same virtual network link, the routing mechanism routes the packet based on the layer 2 destination address specified by the source of the packet.

38. The apparatus of claim 33, wherein the routing mechanism is configured to ensure that the packet includes a hop count to avoid problems with temporary loops.

39. The apparatus of claim 33, wherein the link state determination mechanism is configured to automatically obtain information specifying where endnodes are located by examining data packets.

40. The apparatus of claim 33, wherein the link state determination mechanism is configured to automatically obtain information specifying where endnodes are located by examining control packets.